

WHAT IS CLAIMED IS:

1. A method of reducing a memory footprint of a database table having a plurality of rows and one or more columns, wherein each of the one or more columns has a cardinality, and wherein the cardinality is a total number of different values in the rows of each column, the method comprising:
 - comparing the cardinality with a total number of possible values in the rows of at least one column based on a width of the column; and
 - reducing the width of the column if the cardinality is less than a threshold based on the total number of possible values in the rows of the at least one column.
2. A method in accordance with Claim 1, wherein the threshold relates to a least integer greater than or equal to the logarithm to the base two of the cardinality of the column.
3. A method in accordance with Claim 1, wherein a value of an entry in a row and a column comprises a data entry in a cell, wherein the column in the table has a maximum value length of k bits.
4. A method in accordance with Claim 3, wherein a dictionary for the column has an entry for each different value in the column, wherein the dictionary for the column comprises a width of k bits.
5. A method in accordance with Claim 4, wherein the width of the column comprises a number of bits used to specify column entries, wherein the column comprises a width of w bits, wherein w is an integer, wherein a value of w determines a number of different values in the column, wherein p is the number of different possible entries in the column, and wherein $p=2^w$.
6. A method in accordance with Claim 5, wherein the cardinality of the column comprises a number of different values in the column, wherein the table comprises n rows, wherein the column comprises m different values and has cardinality m, and wherein the value of w is such that $m \leq p$, and if the column has repeated entry values then $m < n$.

7. A method in accordance with Claim 6, wherein $\log_2 p = w$ and $\log_2 m < w$.
8. A method in accordance with Claim 6, further comprising:
rewriting the column with one or more dictionary references; and
resetting the column width to w .
9. A method in accordance with Claim 8, further comprising:
decrementing w in increments of 1 as long as $m < p/2$; and
setting a value of w to w_{\min} when decrementing ceases, wherein w_{\min} is the least integer greater than or equal to $\log_2 m$ for a column with cardinality m and width w .
10. A method in accordance with Claim 1, the method further comprising:
writing a dictionary for the column, wherein the dictionary references the column entries, and wherein the dictionary comprises one row for each of the different values in the column; and
replacing column values by the dictionary references, wherein the dictionary comprises m rows, and wherein each row comprises a width of w bits.
11. A method to reduce an amount of memory associated with information in a database table having a plurality of rows and one or more columns, wherein the information relates to at least two columns, the method comprising:
determining respective values in the at least two columns in a memory;
determining whether the respective values are interdependent;
upon determining an interdependency, generating a combined column based on the at least two columns, wherein the combined column includes the respective values in the at least two columns; and
upon generating the combined column, deleting the at least two columns from memory.

12. A method in accordance with Claim 11 wherein a value identifier (ID) for a value in a dictionary comprises a row number of a corresponding entry in the dictionary, and wherein a document identifier (ID) of a column entry comprises a dictionary reference at a corresponding row number in the table.

13. A method in accordance with Claim 12, wherein the method comprises columns 1 and 2 with n rows and respective document identifiers (IDs) $d1j$ and $d2j$ for $1 \leq j \leq n$, and wherein the method further comprises dictionaries 1 and 2 that are configured to list different values in columns 1 and 2.

14. A method in accordance with Claim 13, wherein the method further comprises a dictionary 12 for combined column 12 that is adapted to list value IDs as pairs $[d1j, d2j]$, for $1 \leq j \leq n$, wherein document IDs $d1j$ and $d2j$ are from row j in columns 1 and 2 respectively.

15. A method in accordance with Claim 14, wherein a document list for the combined column 12 comprises entries $d12j$, for $1 \leq j \leq n$, wherein each entry is adapted to provide the dictionary entry $[d1j, d2j]$ for row j in columns 1 and 2.

16. A method in accordance with Claim 15, the method further comprising deleting columns 1 and 2 from memory upon the existence of the combined column 12, wherein the deletion reduces an amount of memory used to store the data from columns 1 and 2.

17. A method in accordance with Claim 16, wherein column 1 comprises n rows, cardinality $m1$, and a minimum width of $w1$ bits, wherein the dictionary for column 1 comprises length $m1$ and width $w1$, the column 1 dictionary comprising $m1 * k1$ bits in memory, and wherein the document list for column 1 comprises length n and width $w1$, the document list comprising $n * w1$ bits in memory.

18. A method in accordance with Claim 17, wherein column 2 comprises n rows, cardinality $m2$, and a minimum width of $w2$ bits, wherein the dictionary for column 2 comprises length $m2$ and width $w2$, the column 2 dictionary comprising $m2 * k2$ bits in

memory, and wherein the document list for column 2 comprises length n and width w_2 , the document list comprising $n * w_2$ bits in memory.

19. A method in accordance with Claim 18, wherein the dictionary for combined column 12 comprises a maximum length of $(m_1 * m_2)$ and further comprises a maximum of $(m_1 * m_2 * (w_1 + w_2))$ bits in memory.

20. A method in accordance with Claim 19, wherein the document list for combined column 12 comprises a length of n and further comprises a maximum of $(n * (w_1 + w_2))$ bits in memory.

21. A method in accordance with Claim 20, wherein the memory for combined column 12 comprises a measure of an extent of functional dependencies between columns 1 and 2.

22. A method in accordance with Claim 21, wherein the method is adapted to be generalized to combine n columns into a single column.

23. An article comprising a machine-readable medium storing instructions operable to cause a machine to perform operations comprising:

reducing a memory footprint of a database table having a plurality of rows and one or more columns, wherein each of the one or more columns has a cardinality, and wherein the cardinality is a total number of different values in the rows of each column, the reducing comprising:

comparing the cardinality with a total number of possible values in the rows of at least one column based on a width of the column; and

reducing the width of the column if the cardinality is less than a threshold based on the total number of possible values in the rows of the column.